Tools Used in Troubleshooting

This chapter describes the troubleshooting tools available for the Cisco Nexus 1000V.

Commands

You use the CLI from a local console or remotely using a Telnet or Secure Shell (SSH) session. The command-line interface (CLI) provides a command structure similar to the Cisco NX-OS software, with context-sensitive help, **show** commands, multi-user support, and role-based access control.

Each feature has **show** commands that provide information about the feature configuration, status, and performance. Additionally, you can use the following commands for more information:

- **show system—Provides information on system-level components, including cores, errors, and exceptions. Use the show system error-id command to find details on error codes:**

```
n1000v# copy running-config startup-config
[########################################] 100%
2008 Jan 16 09:59:29 zoom %$ VDC-1 %$ %BOOTVAR-2-AUTOCOPY_FAILED: Autocopy of file
/bootflash/n1000-s1-dk9.4.0.0.837.bin.S8 to standby failed, error=0x401e0008

n1000v# show system error-id 0x401e0008
Error Facility: sysmgr
Error Description: request was aborted, standby disk may be full
```

Ping

The ping utility generates a series of *echo* packets to a destination across a TCP/IP internetwork. When the echo packets arrive at the destination, they are rerouted and sent back to the source. Using ping, you can verify connectivity and latency to a particular destination across an IP routed network.

The ping utility allows you to ping a port or end device. By specifying the IPv4 address, you can send a series of frames to a target destination. Once these frames reach the target, they are looped back to the source and a time stamp is taken. Ping helps you to verify the connectivity and latency to the destination.

Traceroute

Use traceroute to do the following:

- Trace the route followed by the data traffic.

- Compute inter-switch (hop-to-hop) latency.

Traceroute identifies the path taken on a hop-by-hop basis and includes a time stamp at each hop in both directions. You can use traceroute to test the connectivity of ports along the path between the generating switch and the switch closest to the destination.

Enter the **traceroute** command to access this feature.

If the destination cannot be reached, the path discovery starts, which traces the path up to the point of the failure.

Monitoring Processes and CPUs

The CLI has features that enable you to monitor switch processes and CPU status and utilization.

Identifying the Processes Running and Their States

Use the **show processes command** to identify the processes that are running and the status of each process. (See Example 2-1.) The command output includes the following:

- PID—Process ID.
- State—Process state.
- PC—Current program counter in hex format.
- Start_cnt—How many times a process has been started (or restarted).
- TTY—Terminal that controls the process. A "-" (hyphen) usually means a daemon that is not running on any particular TTY.
- Process—Name of the process.

Process states are as follows:

- D—Uninterruptible sleep (usually I/O).
- R—Runnable (on run queue).
- S—Sleeping.
- T—Traced or stopped.
- Z—Defunct (zombie) process.
- NR—Not-running.
- ER—Should be running but currently not-running.

Note The ER state typically designates a process that has been restarted too many times, which causes the system to classify it as faulty and disable it.

Example 2-1 show processes Command

```
n1000v# show processes ?
  >       Redirect it to a file
  >>      Redirect it to a file in append mode
  cpu     Show processes CPU Info
  log     Show information about process logs
  memory  Show processes Memory Info
```

```
  vdc     Show processes in vdc
  |       Pipe command output to filter

n1000v# show processes

PID    State  PC        Start_cnt    TTY   Process
-----  -----  --------  -----------  ----  -------------
    1      S  b7f9e468            1     -  init
    2      S         0            1     -  migration/0
    3      S         0            1     -  ksoftirqd/0
    4      S         0            1     -  desched/0
    5      S         0            1     -  migration/1
    6      S         0            1     -  ksoftirqd/1
    7      S         0            1     -  desched/1
    8      S         0            1     -  events/0
    9      S         0            1     -  events/1
   10      S         0            1     -  khelper
   15      S         0            1     -  kthread
   24      S         0            1     -  kacpid
  101      S         0            1     -  kblockd/0
  102      S         0            1     -  kblockd/1
  115      S         0            1     -  khubd
  191      S         0            1     -  pdflush
  192      S         0            1     -  pdflushn
...
```

Displaying CPU Utilization

Enter the **show processes cpu** command to display CPU utilization (see Example 2-2). The command output includes the following:

- Runtime(ms)—CPU time that the process has used, expressed in milliseconds.
- Invoked—Number of times that the process has been invoked.
- uSecs—Microseconds of CPU time as an average for each process invocation.
- 1Sec—CPU utilization as a percentage for the last one second.

Example 2-2 show processes cpu Command

```
n1000v# show processes cpu

PID    Runtime(ms)  Invoked   uSecs  1Sec    Process
-----  -----------  --------  -----  ------  -----------
    1         2754       458   6013    0.0%  init
    2            0       166      4    0.0%  kthreadd
    3            0         2      0    0.0%  migration/0
    4          239     51386      4    0.0%  ksoftirqd/0
    5            2        72     27    0.0%  watchdog/0
    6           12      1798      6    0.0%  events/0
    7            0        27      7    0.0%  khelper
    8           39      2278     17    0.0%  kblockd/0
    9            0         2      0    0.0%  kacpid
   10            0         2      0    0.0%  kacpi_notify
   11            1         9    200    0.0%  kseriod
   12            0         2      0    0.0%  ata/0
   13            0         2      0    0.0%  ata_aux
   14            0         2      0    0.0%  ksuspend_usbd
   15            0         2      1    0.0%  khubd
   16            9       356     27    0.0%  pdflush
```

```
   17            0         5      2    0.0%  pdflush
   18            0         2      1    0.0%  kswapd0
   19            0         2      1    0.0%  aio/0
   20            0         2      1    0.0%  nfsiod
   21            0        19      4    0.0%  rpciod/0
  331           24        63    385    0.0%  kjournald
...
```

Displaying CPU and Memory Information

Enter the **show system resources** command to display system-related CPU and memory statistics (see Example 2-3). The output includes the following:

- The load is defined as the number of running processes. The average reflects the system load over the past 1, 5, and 15 minutes.
- Processes displays the number of processes in the system, and how many processes are actually running when the command is entered.
- CPU states shows the CPU usage percentage in the user mode, kernel mode, and idle time in the last one second.
- Memory usage provides the total memory, used memory, free memory, memory used for buffers, and memory used for cache in kilobytes. Buffers and cache are also included in the used memory statistics.

Example 2-3 show system resources Command

```
n1000v# show system resources
Load average:   1 minute: 0.00   5 minutes: 0.14   15 minutes: 0.16
Processes   :   295 total, 4 running
CPU states  :   0.0% user,   2.0% kernel,   98.0% idle
Memory usage:   2064844K total,   1379800K used,    685044K free
Current memory status: OK
```

RADIUS

The RADIUS protocol is used for the exchange of attributes or credentials between a head-end RADIUS server and a client device. These attributes relate to three classes of services:

- Authentication
- Authorization
- Accounting

Authentication refers to the authentication of users for access to a specific device. You can use RADIUS to manage user accounts for access to a Cisco Nexus 1000V device. When you try to log into a device, the Cisco Nexus 1000V validates you with information from a central RADIUS server.

Authorization refers to the scope of access that you have once you have been authenticated. Assigned roles for users can be stored in a RADIUS server with a list of actual devices that the user should have access to. Once the user has been authenticated, the switch can refer to the RADIUS server to determine the extent of access that the user will have within the switch network.

Accounting refers to the log information that is kept for each management session in a switch. This information can be used to generate reports for troubleshooting purposes and user accountability. Accounting can be implemented locally or remotely (using RADIUS).

This example shows how to display accounting log entries:

```
n1000v# show accounting log
Sun Dec 7 04:02:27 2002:start:/dev/pts/0_1039924947:admin
Sun Dec 04:02:28 2002:stop:/dev/pts/0_1039924947:admin:vsh exited normally
Sun Dec 15 04:02:33 2002:start:/dev/pts/0_1039924953:admin
Sun Dec 15 04:02:34 2002:stop:/dev/pts/0_1039924953:admin:vsh exited normally
Sun Dec 15 05:02:08 2002:start:snmp_1039928528_172.22.95.167:public
Sun Dec 15 05:02:08 2002:update:snmp_1039928528_172.22.95.167:public:Switchname
```

Note The accounting log shows only the beginning and ending (start and stop) times for each session.

Syslog

The system message logging software saves messages in a log file or directs the messages to other devices. This feature provides the following capabilities:

- Logging information for monitoring and troubleshooting.
- Selecting the types of logging information to be captured.
- Selecting the destination of the captured logging information.

The syslog software allows you to store a chronological log of system messages locally or send to a central syslog server. Syslog messages can also be sent to the console for immediate use. These messages can vary in detail depending on the configuration that you choose.

Syslog messages are categorized into seven severity levels from *debug to critical* events. You can limit the severity levels that are reported for specific services within the switch.

Log messages are not saved across system reboots. However, a maximum of 100 log messages with a severity level of critical and below (levels 0, 1, and 2) can logged to a local file or server.

Logging Levels

The Cisco Nexus 1000V supports the following logging levels:

- 0—emergency
- 1—alert
- 2—critical
- 3—error
- 4—warning
- 5—notification
- 6—informational
- 7—debugging

By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. You can specify which system messages should be saved based on the type of facility and the severity level. Messages are time-stamped to enhance real-time debugging and management.

Enabling Logging for Telnet or SSH

System logging messages are sent to the console based on the default or configured logging facility and severity values.

You can disable logging to the console or enable logging to a given Telnet or SSH session as follows:

- To disable console logging, enter the **no logging console** command in global configuration mode.
- To enable logging for Telnet or SSH, enter the **terminal monitor** command in EXEC mode.

Note When logging to a console session that is disabled or enabled, that state is applied to all future console sessions. If you exit and log in again to a new session, the state is preserved. However, when logging to a Telnet or SSH session that is enabled or disabled, that state is applied only to that session. The state is not preserved after you exit the session.

The **no logging console** command that is shown in Example 2-4 disables console logging and is enabled by default.

Example 2-4 no logging console Command

```
n1000v(config)# no logging console
```

The **terminal monitor** command that is shown in Example 2-5 enables logging for Telnet or SSH and is disabled by default.

Example 2-5 terminal monitor Command

```
n1000v# terminal monitor
```

For more information about configuring syslogs, see the *Cisco Nexus 1000V for Microsoft Hyper-V System Management Configuration Guide*.